Refactoring the Curiosity Rover's Sample Handling Architecture on Mars

Vandi Verma, Stephen Kuhn
NASA Jet Propulsion Laboratory,
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
verma@jpl.nasa.gov, skuhn@jpl.nasa.gov

Abstract— **The Curiosity Mars rover sample handling hardware and software were architected assuming that end-to-end sampling operations would occur in a single rover position, from acquisition of a powdered sample with a scoop or drill, through to the cleaning out of all sample residue in the sample chain. However, after analysis of the first drilled samples in Yellowknife Bay, the science team wanted to iterate with additional experiments on Mars and in laboratories on Earth to better understand their results and increase the value of science returned. With the architecture as conceived, the time needed to do so was in direct competition with the exploration of other targets and satisfaction of success criteria during the prime mission. The science team desired the capability to "cache" the sample for future use while continuing progress towards mission objectives by driving away and maintaining use of the robotic arm for contact science. Allowing sample to move about freely in this state risked hardware damage, ending the ability to deliver sample using the nominal path. In this paper we present the approaches that were developed to repurpose some of the sampling hardware into a series of caches and catchments that reduced this hardware risk to a level acceptable during the prime mission. This approach presented new challenges for rover planners, who had to learn to command the robotic arm using new routines that were too complicated to manage without assistance. The rover planner Software Simulation ("SSim") was updated to track the turret gravity vector and sample state, generating an execution error or breakpoint as constraints were violated. Sample from the Cumberland drill target was cached for over 9 months, facilitating a number of scientific discoveries. As data accumulated and the mission transitioned into extended operations, the cached sample capability evolved to significantly simplify operations and reduce overhead.**

TABLE OF CONTENTS

1. INTRODUCTION

The Mars Science Laboratory ("MSL") Curiosity rover landed in Gale Crater in August of 2012 on its mission to explore Mt. Sharp with an envisioned architecture for serial, in situ sampling operations. In part, this reflected the model of every previous Mars surface mission, the instruments of which were inherently designed to observe the immediate environment. It also reflected a focus across the mission of prioritizing resources to ensure that Curiosity could land successfully, with hardware capable of performing its mission; however, given the capabilities of Curiosity's instruments, this architecture was no longer an imperative. Rather, productivity could be increased significantly, and difficult tradeoffs avoided, if Curiosity was able to reconcile its nominal mission with the retention of sample for later delivery to its instruments. This capability came to be known as "caching" sample. It should be noted that this capability is different from one of the goals of the Mars 2020 mission to cache cored drill samples in sealed tube assemblies that it will drop on the surface of Mars for possible return by a later mission.

While not an intentional or designed capability of the hardware, several articulated locations in a mechanism in the robotic arm's turret called Collection and Handling for In situ Martian Rock Analysis ("CHIMRA") were repurposed to control the behavior of sample. Doing so required a chain of custody in the orientation of CHIMRA with respect to gravity, altering the way the robotic arm was commanded to perform its nominal contact science. This paper describes how a new mission capability was conceived, developed and evolved in the midst of regular, safety critical use in flight.

978-1-5386-6854-2/19/$31.00

2. MSL Sampling System Architecture

A detailed description of the sampling chain is outside the scope of this paper and is addressed extensively in [1], [2] and [3] by many of the original designers of the hardware. While it is not possible to usefully describe its actual physical behaviors in this brief introduction, an abstraction of the hardware architecture is necessary to understand the utility of the capabilities that would be developed to repurpose that hardware.

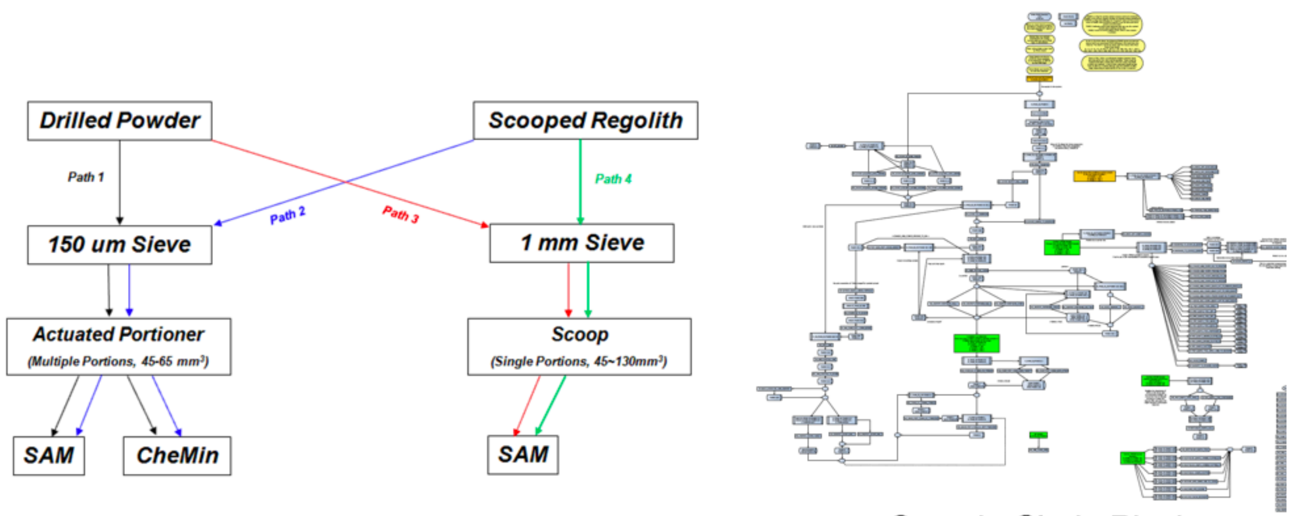

Sample Chain Block Diagram at Sol 0

Sample Chain Block Diagram at Sol 1000

Figure 1. Comparison of the Sample Chain at Landing and Three Years Into Mission

MSL sampling hardware and software using the primary 150µm sampling path were designed with a fixed expected chain of operations: (1) acquire powdered sample with the scoop or drill; (2) sieve that sample to particles of 150µm or less in at least two dimensions; (3) portion that fine sample into a controlled volume; (4) deliver that portion to a cup or holding cell in either of the two sample analysis instruments in the rover chassis; (5) repeat steps 3 and 4 as desired (e.g. to drop off to the other instrument so that both may consider the same sample); (6) dump the residual sample; and (7) clean the sampling hardware to reduce cross-contamination of the next sample.

Explicit in this original architecture was an expectation that these steps would be performed serially and in order, and that no other robotic arm activity would interrupt the flow of sample in this chain. Another sample path using a 1mm sieve and different physical hardware was even more limited, by its inability to prepare multiple portions from the same sample. It was used only twice, and is not discussed here.

This paper addresses one of several vectors for the evolution of the sample chain over the course of the mission. For purposes of illustrating the extent of its evolution, Figure 1 (intentionally illegible due to restrictions on disseminating command dictionary content) portrays the extent of change in the structure of the sample chain from the time of landing to a time of peak mission productivity during the Pahrump Hills walkabout.

3. The Architectural Concern

What will the sample do?

The interaction of Martian sample with the drill and CHIMRA sampling and instrument hardware that would interact with it was a principal research focus in terrestrial characterization of the sample chain. Hundreds of tests were performed in qualification test chambers on a range of plausible terrestrial rock and regolith analogues, at representative atmospheric pressure, humidity and temperature. However, these test chambers could not fully capture the electrostatic conditions that prevail in the Martian atmosphere, nor anticipate the precise composition of the rocks that would be encountered. Furthermore, some samples did behave in ways that seemed to justify the significant expense in time, money and mass that were devoted to mitigations for sample that refused to flow freely, or worse, flowed freely for a time, but became sticky due to some chemical evolution or the physical stimulus of vibration.

Mitigations for "Sticky" Sample

The study and interaction of sample on Mars was devoted a "sample playground" on the rover that included an observation tray, cut-away instrument dropoff funnel, scratch post, and prong to "poke" a plug loose from the portioning part of CHIMRA. However, the greatest concern was blinding of CHIMRA's 150µm sieve. In every other part of the nominal 150µm sample chain, sample flowed through passages that were at least several square centimeters in size, until the final portion was routed into a tube and then dropped into a funnel narrowing to a few mms in diameter. However, every bit of sample that would be portioned and dropped off to the instruments had to pass through 150µm-wide etched perforations in a 50µm-thick titanium sieve. The Chemistry and Mineralogy X-ray Diffraction ("CheMin") instrument required particles this small and smaller, and although the Sample Analysis at Mars ("SAM") instrument could accommodate larger sizes, the architecture could not support independent paths for each instrument to deliver the same sample.

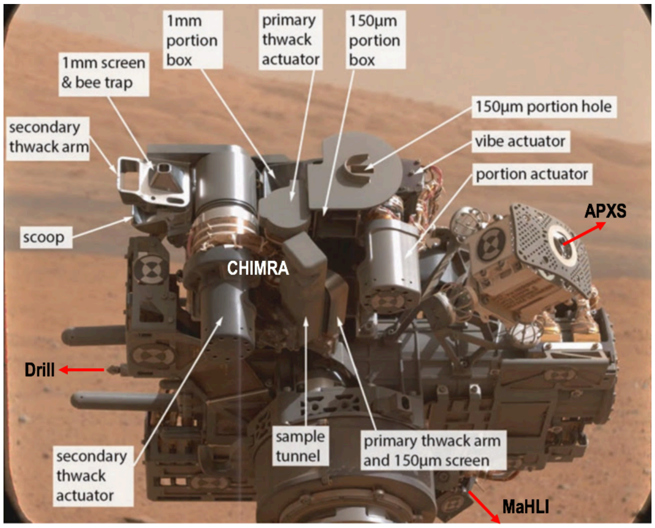

Figure 2. Annotation of CHIMRA and the APXS and MaHLI Instruments Clocked About it on the Turret

Because blinding of CHIMRA's 150µm sieve was the greatest passive hardware safety concern in the sampling system, the designers devoted more or less an entire actuator to mitigating this risk. While it serves some diagnostic utility in exposing parts of CHIMRA, and also allows sample flow through the portion cone, or tube, to be minimized, it was not technically required in the architecture, except to perform the "thwack."

The thwack mechanism uses the thwack actuator to rotate the piece of CHIMRA hardware (called the "tunnel," upon which the sieve is welded) away from the chamber to which it mates when sieving (the "150µm reservoir"). A torsion spring stores that energy until, at a fixed position, a latch and tang mechanism releases, and the tunnel is accelerated towards the 150µm reservoir by the force stored in the spring. The acceleration into the tunnel is substantial, but the thwack at impact between these two pieces of hardware imparts some of the highest dynamic forces any robotic mission ever generates intentionally – measured to be more than 5,000g. In the terrestrial testbed, it scares the dickens out of unsuspecting observers. Critically, the point of this mechanism was to overwhelm whatever forces had contributed to getting the particle stuck.

The direction of the thwack force expels particles that became clogged while sieving, out in the direction from which they came. The sieve was installed in such a way that the Venturi bevel (simplified for purposes of illustration in Figure 3) inherent to the etching process that created the sieve perforations was oriented towards the post-sieved side, where it would limit the contact area with which a particle could become press fit into a narrowing opening while sieving.

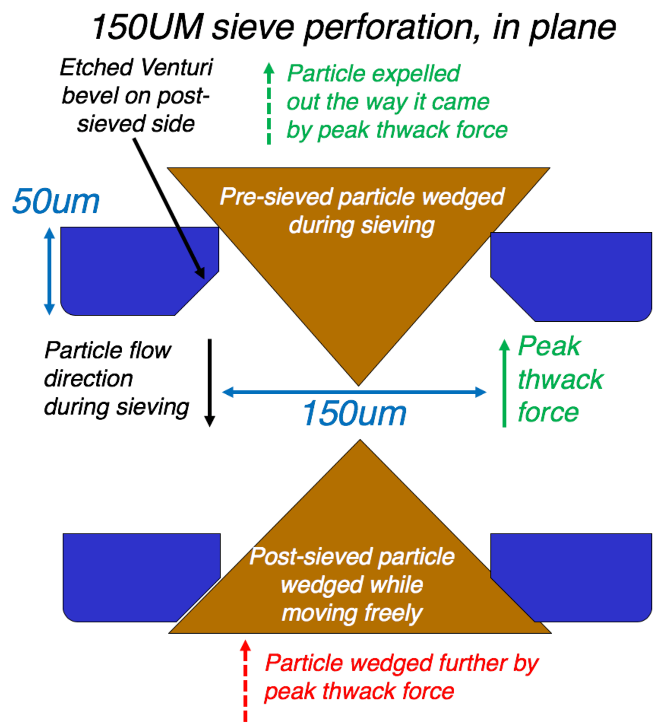

Figure 3. Use of Thwack Potentially Ineffective/Counter-Productive in Clearing Particles Clogged on Downstream Side of Sieve

Continuity of Mitigation

Early in the mission, before a range of samples had been acquired, and with prime mission objectives yet to be completed, the engineering team had not altered the risk posture it had maintained throughout development.

In considering the viability of caching and managing the cached sample during other rover activities, one foundational constraint was given deference. If sample was allowed to interact with the reverse side of the sieve, attempting to clean particles that had become lodged in the sieve using the thwack would exert peak forces that would merely wedge the sample further. The interaction could never be overwhelmed with greater forces, because the thwack was that greatest force. Therefore, sample that had already been sieved must not be allowed to interact with more than a fraction of the sieve.

4. ORIGINS OF CACHED SAMPLE

The desires for caching sample were two-fold.

Engineering Logistics

First, there was a logistical concern with the phasing of operations oriented around the sample. Once a sample was dropped to the observation tray, the Mars Hand Lens Imager ("MaHLI") on the turret (see Figure 2) provided the best quality images, and the Alpha X-ray Particle Spectrometer ("APXS") instrument provided the only plausible way of pre-assessing its makeup. To use these contact science instruments, however, the robotic arm must command them into position along the target's axis.

This was simply not a part of the seven-step, serial architecture previously described. In placing these instruments, the sample was not moving on its pre-ordained path through CHIMRA's chambers to the portion hole from which it would be delivered to instruments. Rather, it was just sloshing around. Though spilling the sample or backflow through the 150µm sieve were not concerns, interaction with the reverse side of the sieve certainly was. Engineers also weren't excited about packing sample into the portioning tube without immediately dropping it off to an instrument.

Having it Both Ways

More important to the mission timeline than these capabilities to assess the sample in situ before dropping it off to the instruments was the ability to drive away from the site at which the sample was collected and continue performing nominal contact science.

The second drilled sample Curiosity acquired, at Cumberland in Yellowknife Bay, was cached from Sols 282 to 487. This sample would prove to be exciting for its high concentration of organics, above levels believed to be attributable to error and contamination. [4] During that time, Curiosity performed contact science campaigns at Point Lake, Shaler, Darwin (Waypoint 1 on the transit from Yellowknife Bay to Mt. Sharp) and Cooperstown (Waypoint 2). During that time, the Sample Analysis at Mars ("SAM") instrument received eight separate Cumberland portions (often a SAM "portion" was comprised of three separate CHIMRA drop-offs, as SAM sought to isolate how results varied with sample mass and improve their signal to noise) that it was able to analyze with different parameters, tuning subsequent experiments to reflect the results of the previous. While SAM is also able to reuse its analysis cups until residue completely fills them and to store delivered samples indefinitely (so called "doggie-bagging"), the pristineness of a cup is a sort of soft consumable. SAM also has different types of cups, including precious wet chemistry cups that were not reusable.

Caching the sample allowed scientists to reconsider their desires after each sample analysis, rather than estimating ahead of time (and on the clock) how many different analyses they wanted to perform. A truly exciting sample, like Cumberland, merited numerous experiments representing significant expenditure of energy, and of pristine cups. However, every one of the dozens of drill samples that were expected at the start of the mission to eventually be acquired could not be so lucky. It took more than an hour to prepare and deliver three CHIMRA aliquots to a single SAM cup. SAM could not "dump" its cups, but could only make them ready for reuse with the energy-intensive combustion of the analysis itself, which consumes nearly all of the discretionary energy for a Sol. Therefore, "contingency" doggie-bagging was not a sustainable proposition without impacting mission cadence at some point in the future. Instead of devoting time and energy to sample that could never be analyzed, the SAM team was able to deliberate in consideration of their initial analysis results and take only the samples they needed.

5. PREMISES OF CACHED SAMPLE

Two core observations informed the principles of cached sample.

The Angle of Repose

First, powdered sample at rest does not behave like an ideal fluid but can be relied upon to hold an angle of repose. A principle familiar to civil engineers in their prescriptive requirements for fill slopes is even more pronounced in sample with an electrostatic tendency to be sticky, as fine, roughly angular powder does, especially when vibrated. The "angle of repose" was the angle with respect to gravity that a cohesive sample could maintain without breaking its shape.

In the context of sampling on Curiosity, this principle applies to the way sample is generally moved from place to place in a controlled manner. A given turret pose is commanded using the five degrees of freedom of the arm, and the CHIMRA vibration actuator (a motor shaft with an eccentric mass attached) spins for a number of seconds. Then, a change to the turret gravity vector is commanded that does not exceed the angle of repose of that sample, and the process repeats. During vibration, finely sieved sample on a modest slope flows like warm syrup, seeking out the low places without getting too enthusiastic about it and overshooting. After vibrating, the sample tends to form a plane. Figure 4 shows an image of sample that was vibrated about 30 degrees from the pose it struck in the image, into the back of the portion box.

The Catchments

Second, CHIMRA was designed to perform a number of functions in a tight volume with as few actuators as possible. This tended to bias its design towards a physical form in which a progression of orientations and shared degrees of freedom takes the sample on a winding path. It is not really a labyrinth, but it does have a number of good hiding spots behind partitions that perform functions in the nominal sample chain.

By placing sample on one side of these partitions, they could serve as catchments to separate the sample from contact with the sieve. Up to a point.

The Strategy

Essentially, the strategy of caching sample was to find a set of catchments that could hold sample at different orientations of the turret. Typically, a given instrument would have two to three different catchment locations that would enable it to be pointed in every conceivable attitude for contact science – all of the poses within 90 degrees of pointing straight down. Figure 5 portrays this for the MaHLI instrument. The false

Figure 4. Sieved Sample Holding an Angle of Repose on Mars in the CHIMRA Portion Box

color images portray powdered sample (in pink) in the CHIMRA portion box (yellow) and tunnel (blue) to which the sieve (light orange) is mated when the tunnel is closed. To the left of each image, the rover is portrayed with the turret in the orientation that corresponds to that sample state. In some orientations, sample would be expected to spill (dotted red lines and arrows indicating sample breaking repose). Limited contact with the sieve (e.g., upper left, with the contact surface highlighted in green) was permitted.

The criteria for a catchment location was that when the maximum supported sample volume (about 12 cubic centimeters) was vibrated into position, none would contact the reverse side of the sieve. From each of these positions, a range of movement could be tolerated. This could correspond to holding an angle of repose (see two upper right images in Figure 4) or, alternately, spillage to a location that results in contact with no more than a small fraction of the sieve (remaining images in Figure 4). But once a sample is believed to break repose, it must be "prepared" anew before it is allowed to slosh any further. And before it can be prepared, any contact with the sieve must be "recovered" as best as possible by safely attaining the most favorable pose for expulsion, and vibrating there. Because vibration forces are no more than an order of magnitude greater than a Martian g,

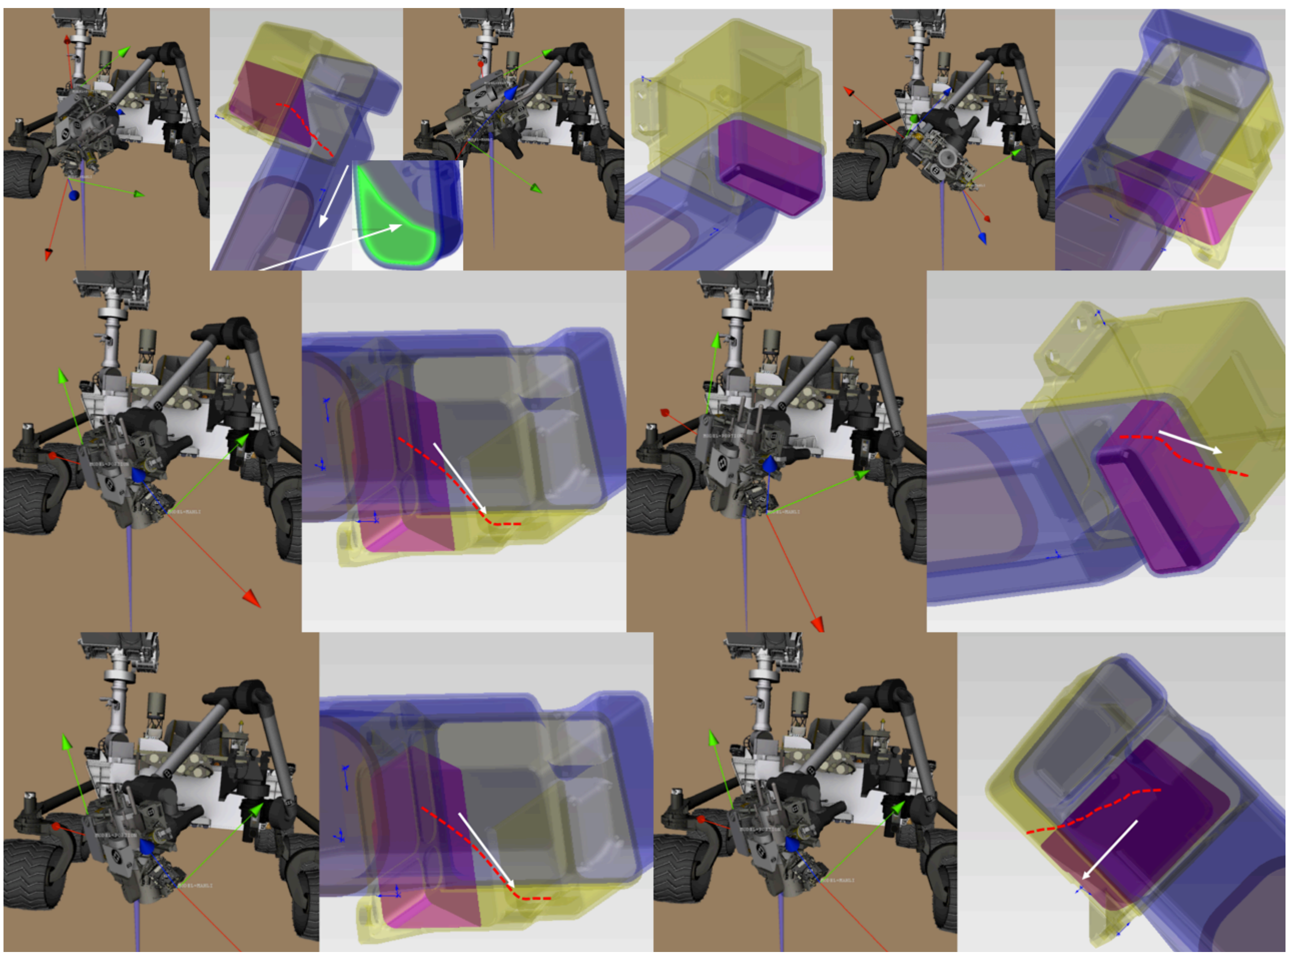

Figure 5. Cache Locations for the MaHLI Instrument in the Seven Zones Created Within 90 Degrees of Tool Alignment With Gravity.

and irregular, this was not believed to be a panacea for potential damage, merely a best practice.

6. The Approach

The Need to Abstract Sample State

Even for the engineers responsible for CHIMRA, mentally tracking sample across a range of orientations as it broke repose and fell through the various partitions in CHIMRA was a losing proposition. If cached sample was to be viable, rover planners needed a way to know whether what they were doing was permitted. While this sort of checking could be implemented in the Rover Compute Element ("RCE") Flight Software ("FSW") where many other high-level behaviors reside, the bar for updates to FSW was high, with long latency. Since landing in 2012, neglecting patches and other scripting workarounds, Curiosity's FSW has only been updated to a new version twice. Furthermore, it was clear that the capability was a moving target, that updates would be made, and that they would have to be timely.

Autonomy with SSim

A key capability of the rover planner's ground tools is the ability to quickly and repeatedly simulate their commands using the actual code that would execute on the spacecraft. Surface Simulation ("SSim") [5] provides this interface. At any point throughout the development of their activities, rover planners are able to use SSim to simulate in seconds what would take hours to execute in real time on the rover.

SSim simulates flight software by using actual flight software code. As a result, it can query FSW state even if that state is not exposed via telemetry and commands. SSim can also be programmed to augment the behaviors of Curiosity's RCE FSW. Because SSim is a part of the suite of ground tools used on MSL, it can be updated with a much nimbler process than the FSW. Therefore, this is where the cached sample monitoring behavior was implemented.

To execute a single robotic arm motion command, FSW generates a trajectory that may consist of dozens of via points. The trajectory is then executed through a closed loop motor control request. Since a cached sample violation could occur anywhere along this motion trajectory, it was not sufficient to check violations only at the start and end of a command, or even only at the via points. SSim monitoring would be continuous.

An existing framework of commanded variable setting as shorthand for sample states was augmented to coordinate this monitoring. The sampling team had long used three commanded variables embedded in sample processing scripts in order to track what had become a network of relationships in the sample chain and assert if something was attempted that was unintentional. (See the structure encoded in Figure 1 at right). One variable encodes the nodes in the graph, which was structured such that any constraint checking could use either simple equality or a range of valid values (i.e. be expressed with only numeric comparison operators). Another variable encoded the concept of transiency, that a behavior that faulted in progress would need to be completed manually before moving on. A third variable communicated whether sample was present on either side of the 150µm sieve or on both sides.

To this were added two other variables tracking cached sample. A binary variable acted as an enable for the SSim constraint checking, which typically began at the end of the preparation of a sample into a catchment. Another variable would always be set immediately prior, communicating to SSim which catchment location was being used and therefore which orientations were permitted. In addition to CHIMRA cached sample monitoring, the same variables also tracked sample in the Drill bit chambers so that it would not spill from the Drill Bit Assembly ("DBA"). The DBA is nominally used to transfer sample from the Drill to CHIMRA.

After preparing the sample, rover planners would command their contact science while SSim monitored the turret orientation. As a consequence, unlike the FSW, which modeled an arm movement in its entirety before attempting, and which refused to begin the movement if it could not be completed successfully, SSim would effectively create an execution breakpoint at the precise location in which a constraint was violated. It would also generate its own warning messages describing the issue, interleaved and in proper context with the "real" event reports from the FSW. This was extremely useful in diagnosing what had gone wrong, especially near the triple points of the plot where it is difficult to eyeball attitude. Five types of requirements were encoded in SSim.

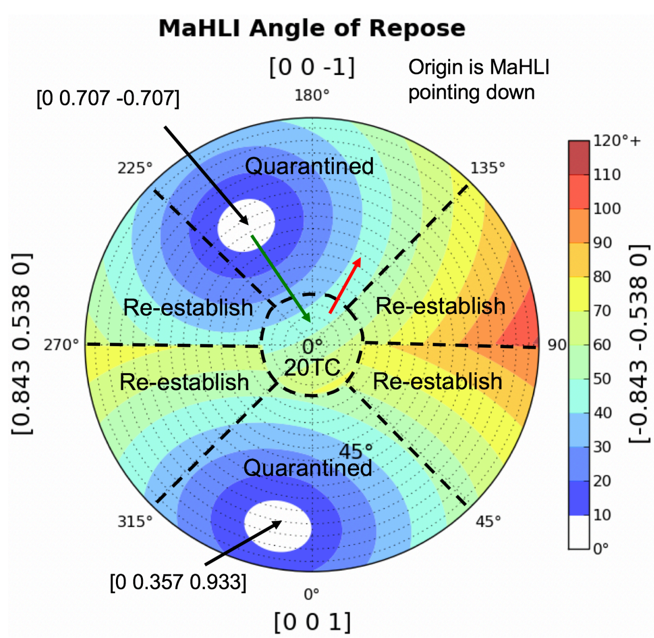

Figure 6. The Seven Different Zones of Turret Orientation for the Two Complementary Strategies for Caching Sample When Using the MaHLI instrument

First, any movement into the zones marked "re-establish" or "20TC," short for 20-degree tilt cone, in Figure 6 had to be performed with a single command from the preparation location (the centers of the white ovals in Figure 6). Typically, this would be a simple turret and wrist move in the arm pose in which the sample was prepared. This ensured that changes in orientation did not take a winding path. As portrayed in Figure 6, the ultimate desired orientation would control which of the preparation locations were used, in order to reasonably minimize the angle of repose.

Second, once the single move into the "re-establish" zone, was complete, no deviations in orientation were tolerated outside of a small halo that permitted the minor perturbations that result from discretizing waypoints when commanding orientation-preserving robotic arm movement. Here, the sample would be expected to have spilled, and it ought not be sloshed around.

Third, once the turret orientation had transitioned to the 20TC, it could not leave that range of orientations before recovering the sample. Here, the angle of repose to avoid spillage could be as high as 65 degrees, and if sample broke, it would contact the sieve. Because this was the most heavily used zone (corresponding to flatter ground and more vertical target normal), allowing some range of travel was more important than in the re-establish zones. However, limiting the range of travel kept the potential damage to a sustainable level.

Fourth, it was permitted to move freely within the "quarantined" zone in which the sample was prepared, but once the orientation had moved from the point of preparation, it was not permitted to enter the 20TC or any re-establish zone.

Fifth, constraints were placed on the use of the other CHIMRA actuators while sample was cached. For instance, it was not permitted to vibrate with sample cached because this would defeat the purpose of squirreling away the sample at some angle of repose.

SSim tracked whether sample had been prepared, the preparation strategy, and the remaining permissible behavior with an internal SSim variable with 34 possible values. SSim had additional state variables to track commands that only change the turret orientation since only a single linear change in orientation was allowed to transition into some zones. In addition, SSim monitored constraints on CHIMRA tunnel, portioner and vibration actuators.

When the range of permissible orientations had been played out, a recovery script was executed to reset the SSim constraint checking by setting the enable variable to disabled. At this point, the sample completed its excursion and was back on a known node in the sample chain. From this state, additional nominal sample chain operations or a new preparation of cached sample were possible.

Keeping these requirements straight as the robotic arm moved about the workspace was simply not something that any rover planner could do. That the ground tools had been architected with SSim's capabilities was a sort of foresight and investment that transcended the conception of cached sample.

Learning a New Way to Command

From Sol 65, when the first iterations of cached sample began, to approximately Sol 950 in the Pahrump Hills when cached sample evolved to a second phase of innocuousness, rover planners sequenced contact science with cached sample using techniques that deviated significantly from those without. Essentially gone was the useful command aggregating the behavior of a number of more primitive arm commands that however do not preserve the turret orientation.

Instead, after the sample was prepared, the turret orientation had to be kept within a narrow range. Commanding the arm to move the turret while maintaining its orientation with respect to gravity introduces kinematic constraints that make certain behaviors difficult. First, because the robotic arm possesses only five degrees of freedom, the full set of physically possible turret orientations with respect to gravity can only be accessed from a single shoulder position that varies across 180 degrees depending on rover tilt. Because both preparation and recovery cached sample processing behaviors access turret poses that are only kinematically possible above the chassis deck in these shoulder positions, switching to a new target of sufficiently different orientation required ascending from the ground.

Large moves that preserve the turret gravity vector are also prone to transitioning into kinematic requirements that can't all be reconciled. For instance, one arm command used with cached sample must move through Cartesian space in a straight line while preserving the turret orientation with respect to gravity. Where either of these constraints can't be satisfied kinematically, a trajectory generation or joint limit error will be thrown by the robotic arm FSW, and the command will fail before any movement occurs. Trying to internalize the kinematics of the arm to quickly identify why such an error would be thrown and how to fix it was a skill that took time and practice to develop.

Command Help

Rover planners on MSL are required to author sequences containing hundreds of commands with thousands of arguments in the space of a few hours. It would be impossible to do this without attempting to standardize some of this work into patterns encoded in "macros." A macro could output a standard set of commands to move MaHLI with respect to a target that was defined outside of the macro, in a "standard suite" that includes a 25cm context image, 5cm stereo images, and a 1-2cm detail image. A macro could also command a single portion preparation and drop it off to SAM at the height it calculated could be supported at the current tilt. A macro could also come to do some of what was required to cache sample, finding the desired sample preparation method for the target name and tool that were input.

The complexity and work to implement macros tends to increase as the permutations of inputs expand. The cached sample macro was one of the more complicated and frequently had to be executed several times in a planning cycle for different targets or for different tools on the same target. The robotic arm frequently uses one of three different kinematic arm configurations for contact science, and the macro supported transitions from any one to any other for nearly every catchment strategy. It supported use of MaHLI, APXS and also the Dust Removal Tool ("DRT") across their 90-degree orientation cones with respect to gravity (which together spanned every possible turret orientation) while sample was cached. It also provided output to help a user tweak the orientation of a target for more efficient merging near the boundaries between management zones.

7. THE SCIENCE THAT WAS BOUGHT

Through Sol 1540 when the MSL drill feed experienced its mission-altering anomaly, MSL performed contact science on approximately 836 separate targets (aggregating rasters as a single target). This corresponded to approximately 403 unique APXS integration positions and 3,587 unique MaHLI image positions (at different target offsets and/or raster locations around a single target).

Approximately 404 of those targets were encountered while sample was cached in CHIMRA. This corresponds to 201 unique APXS positions and 1,681 unique MaHLI positions, or very nearly half of the contact science performed. While this productivity was not fully attributable to the ability to cache sample (since a different balance of instrument analysis would certainly have been struck if the mission lacked the capability), many of these targets would not have been possible without it.

Cached sample became arguably more important when the mission hit its stride during the walkabout at Pahrump Hills and in Marias Pass from Sols 755 to 1143. By most metrics, this has been Curiosity's most sustainably productive period of sampling and contact science. This period contained six of the drill samples Curiosity has successfully collected to date, including Confidence Hills, Mojave, Telegraph Peak, Buckskin, Big Sky and Greenhorn. Nearly 1/3 of all the contact science the mission has performed occurred during this year from late 2014 to late 2015, a period of intense activity for rover planners. While the percentage of contact science performed with sample cached was somewhat lower than over the mission overall, at about 40%, the phasing of its use enabled quick turnarounds from one target to the next. Indeed, it was common to dump the sample from the previous drill acquisition as the rover was reconnoitering the next one.

8. THE PRICE THAT WAS PAID

More Work For the Same Productivity

Before Sol 950, on Sols for which rover planners were only performing arm activity (which tend to be more dense than those that also include a drive), rover planner sequences contained an average of 30 non-reusable movement commands per target encountered with sample cached, compared to 22 without. In other words, the number of commands upon which rover planners exerted most of their intellectual capacity – the motion commands that were written anew for that Sol and not part of some "meta-command" implemented with a reusable sequence – averaged nearly 40% greater. This overhead encompassed the commands to prepare and recover the sample, to ascend and descend from the poses where this occurred, the greater primitivity of these commands compared to those used without caching sample, and other considerations.

While this is a significant increase in and of itself on a tactical timeline, what is not communicated in this statistic of added command overhead is the difficulty of doing all of the things that go into contact science – the target adjustment, changes in arm configuration, traverses and ascents – when the progression of turret gravity vectors is no longer allowed to move freely.

While productivity did not decrease substantially, some drop in throughput compensated to unburden rover planners some of the time. Before Sol 950, on Sols for which rover planners were only performing arm activity, the number of unique MaHLI image positions per Sol of execution (the contact science productivity metric with perhaps the highest signal to noise) was approximately 6.1 with sample cached, compared to 8.8 without. However, the number of unique targets per Sol of execution stayed as high or higher with sample cached as without.

9. THE EVOLUTION OF CACHED SAMPLE

Like most other aspects of the mission, the use of cached sample evolved from its initial formulations. However, improvements to the formalized sample management zones, SSim monitoring, command macros, documentation and training were incremental changes. In the midst of the campaigns at Pahrump Hills, the architecture of cached sample itself would be refactored to render it transparent for nearly every aspect of contact science.

To do so required casting aside the foundational constraint of the behaviors that had been executing for nearly three years. It would become permissible to allow post-sieved sample to move freely and to contact every part of the sieve.

Allowing sample to move freely meant that the robotic arm could be commanded through any series of orientations without modeling the state of the sample vicariously with turret attitude. The joint space and other commands that were typical when no sample was present could be used. In fact, commanding would be no different at all than when no sample was cached in the system, with the minor exception of APXS placement to contact on sand or dumped sample, for which a vibration cleaning operation was still required and demanded some management.

Changing Risk Posture

Deciding when and how to justify useful behavior that adds meaningful risk is a difficult engineering judgment. The end of the prime mission, and with it the anticipation of the sorts of failures that may (and indeed did) lurk just beyond tested design lifetimes, is a milestone that is well-suited, if not always sufficient, to justify a new risk posture.

As Curiosity embarked on its Pahrump Hills campaigns and began acquiring drill samples beyond the two collected in Yellowknife Bay, a body of evidence came into existence that Mars samples were behaving much more benignly than had been feared. Alone, such anecdotal results were little more than conclusory. However, the lack of proximity to any real concern in how the samples collected were behaving formed one leg upon which the introduction of greater risk could stand.

Serendipity With Another Anomaly

Sometimes an engineering anomaly will bring with it a silver lining.

In terrestrial testing on Earth shortly after landing, an engineering model of CHIMRA exhibited a hardware failure. Along the edges of the sieve where it was spot-welded to the tunnel (see dimpling around edge of sieve in Figure 8), some of the welds had popped. As the obligatory Anomaly Response Team ("ART") spun up and characterized the progression of the failure, it was clear that the unzipping could progress to a point that the sieve could deform, ceasing to form a seal and also potentially risking the ability to close the tunnel, which would end the capability to portion and retain sample.

The engineering team developed and began acquiring a series of diagnostic images to help resolve whether this failure mode was incipient on the flight vehicle.

Some required particular incidences of light. For instance, the vignette portrayed in Figure 7 reconciled two constraints. First, the turret had to be placed so that the higher resolution right eye of the Mast Camera ("MastCam") instrument could resolve the plane of the 150μm sieve through a gap between the APXS contact plate (left) and the CHIMRA portioner actuator rock guard (right). Second, light had to be incident at such an angle that a sliver of the tunnel behind the sieve would be visible and not shadowed by the tunnel or the thwack arm, within a tolerance of about 5 degrees. The passage and consistency of light through the sieve demonstrated that it was not blinded. Another technique is portrayed in Figure 8, for which the lighting tolerances were even tighter. This technique required that the sieve again be visible to the right eye of the MastCam. However, sunlight had to be reflected off of the sieve and directly into MastCam to induce a specular reflection.

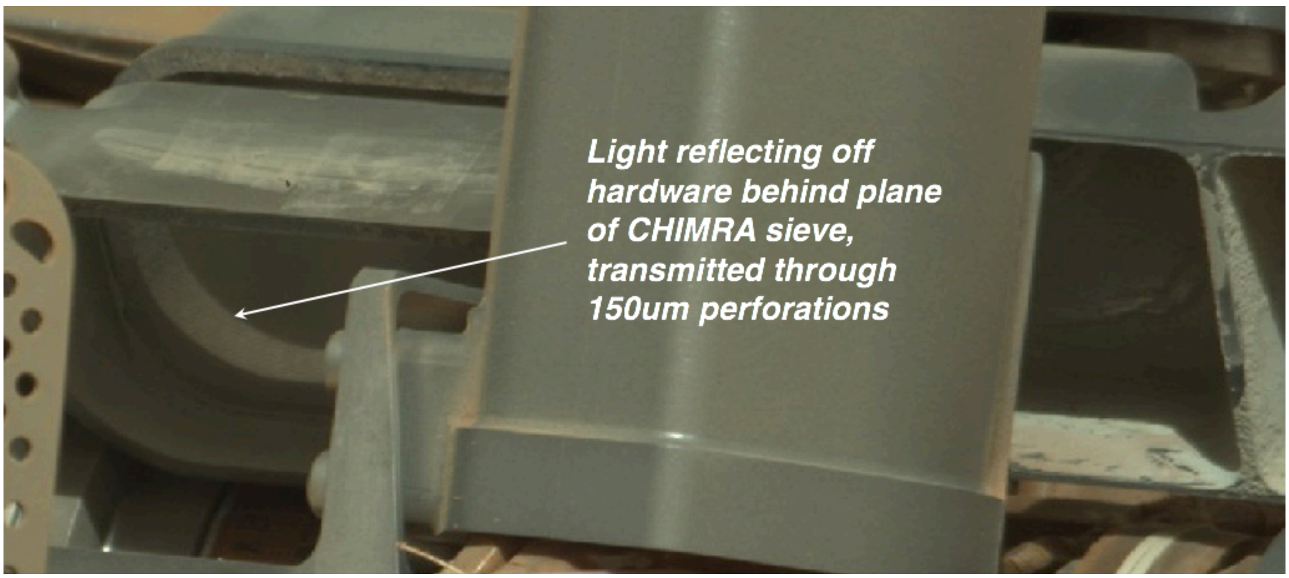

Figure 7. One of the Diagnostic Image Techniques to Help Track the Health of the 150μm Sieve

However, the most useful technique for purposes of risk reduction that the sieve would be degraded by blinding used the ChemCam Remote Mast Imager ("RMI"). This camera has the highest resolution on the rover, and when the turret was positioned near its minimum focus distance, an image could be focus-stacked and stitched that resolved the sieve in high definition. In fact, every one of the tens of thousands of holes in the sieve could be individually resolved at a resolution of about 5 pixels across, providing a basis to track the health of the sieve across samples. Figure 9 is an example of such an image, which also demonstrates the ability to resolve a clog. The ability to trend the health of the sieve with high confidence provided another leg upon which to stand.

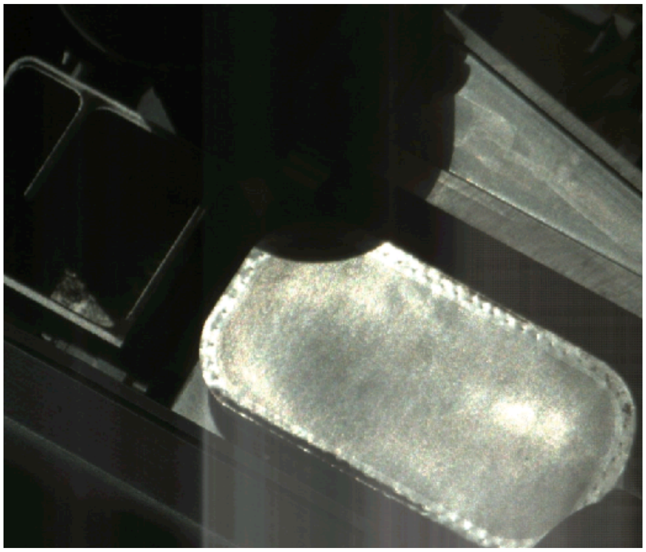

Figure 8. Specular Reflection Off of 150μm Sieve

The Analytical Underpinnings

The sieve de-weldment anomaly brought with it not only flight diagnostics but a host of instrumentation that had not previously been conceived or justified. Among these were high-speed video of the act of thwacking. In one of these tests, sample was introduced. Sieved particles that had become clogged "shed," or were left behind, when the tunnel accelerated towards the impact of thwack, acting to remove sample from the sieve. While this acceleration is two orders of magnitude lower than that experienced during the thwack impact itself, it was still 10-20g depending on the distance from the axis of rotation, significantly greater than the dynamics available with vibration. While not quite as overwhelming a force as the thwack, the empirical evidence of this phenomenon provided the third leg.

Nearly Two Years of Increased Productivity

The second half of the campaigns at Pahrump Hills, including nearly all of the contact science at Marias Pass, were performed with this evolved formulation of cached sample, which was quickly implemented after approvals.

During the more than 500 Sols that this second phase of cached sample was in use before a flight anomaly ended the sample chain as it had been architected, approximately 346 unique targets were sampled, corresponding to approximately 165 unique APXS integration positions and approximately 1,619 unique MaHLI image positions (as Marias Pass became the locale where the art of the massive MaHLI mosaic was most in vogue).

Of these, sample was present in CHIMRA for approximately 248 unique targets corresponding to approximately 119 unique APXS positions and approximately 1,180 unique MaHLI image positions, or roughly three quarters of all contact science. With little price to pay and having achieved an efficient cadence of dumping sample as preparation for immediate acquisition of the next one, the increase was logical. Of those Sols with sample cached for which rover planners performed only contact science and not mobility, the number of non-reusable movement commands per target dropped back down from the average of 30 with the original formulation of cached sample to less than 20 with its

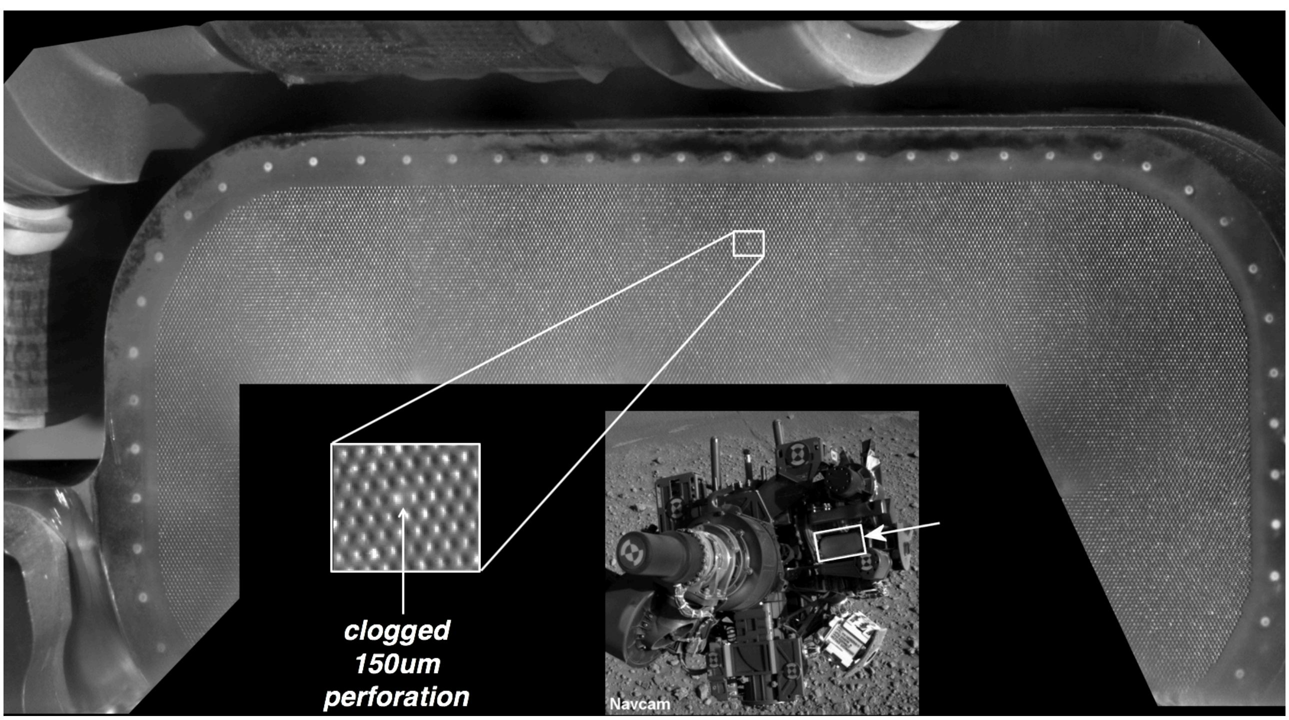

Figure 9. High Resolution Imaging to Help Track the Health of the 150μm Sieve

evolution. The average number of unique MaHLI image positions per Sol of execution returned to 8.9.

10. The Impact of the Drill feed Anomaly

In the extended mission, more than 1,500 Sols after landing, the drill feed mechanism exhibited anomalous behavior that called into question the ability to continue moving the feed at all. After months of diagnostics, it was decided to fully extend the feed such that the drill bit was out as far as it could be placed and leave it there. This created a spatial relationship that enabled the arm to apply force on the bit and actually drill a hole without first engaging the fixed prongs to either side of the drill. However, the permanent extension of the feed ended the ability to transfer sample from the drill to CHIMRA, where all the processing and portioning steps of the sample chain had been performed.

It is the sort of loss of capability that would be conceived as a single-point failure in traditional failure mode criticality analyses. Nevertheless, the spirit of remote robotic operations is to embrace the workaround. Capabilities that came to be known as Feed-Extended Drilling and Feed-Extended Sample Transfer repurposed the drill into both collection and delivery mechanism. It is possible to place the drill bit over the instruments and regurgitate a portion back through the drill augur from which the sample was first acquired, applying a sort of empirical, open-loop portioning with timed vibration from the drill percussion and CHIMRA vibration actuators.

Repurposing the hardware in this way continues a venerable tradition at JPL and within NASA of engineering adaptation until ingenuity is exhausted.

11. Summary

Though not an intentional capability of the Curiosity sampling hardware, procedures were developed to cache sample in ways that permitted a parallelization of mission objectives. Sampling engineers sustainably evolved the capability to maintain and enhance productivity during critical contact science campaigns. The ability to implement complex autonomous constraint checking in SSim was indispensable to changes in behavior that could not have been anticipated when the surface mission began. Future robotics missions may continue to benefit from permitting flexibility in repurposing a fixed hardware design to accommodate changing needs and desires in the pursuit of science. Where this flexibility exceeds the capacities of operators to safely oversee, it may only be viable where adaptation is anticipated and existing, validated methods may be leveraged. Otherwise, the delay, risk and cost associated with coloring outside the lines of an existing test program, when coupled with the many other inherent challenges of robotic operations, provide inertia that often cannot be overcome.

Acknowledgements

The authors thank the many engineers at JPL that have collaborated together to re-architect the sampling system with cached sample, a partial list of whom include Cambria Logan, Joseph Carsten, Daniel Limonadi, Louise Jandura, Chris Roumeliotis, John Michael Morookian and the MSL rover planners.

The authors would also like to thank the Principal Investigators for the ChemCam and MastCam instruments, Roger Wiens and Mike Malin, who along with their teams generously supported the use of their instruments for engineering purposes for many years. Thanks to Olivier Gasnault, who helped focus-stack and stitch a large number of RMI images of the rover hardware over the years.

The work described herein was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

[1] R. C. Anderson, L. Jandura, A. B. Okon, D. Sunshine, C. Roumeliotis, L. W. Beegle, J. Hurowitz, B. Kennedy, D. Limonadi, S. McCloskey, M. Robinson, C. Seybold, K. Brown, Collecting samples in Gale Crater, Mars; an overview of the Mars Science Laboratory Sample Acquisition, Sample Processing and Handling System. Space Sci. Rev. 170, 57–75 (2012). doi:10.1007/s11214-012-9898-9

[2] D. Sunshine, Mars science laboratory CHIMRA: a device for processing powdered Martian samples, in Proceedings of the 40th Aerospace Mechanisms Symposium (2010)

[3] A. Okon, Mars science laboratory drill, in Proceedings of the 40th Aerospace Mechanisms Symposium (2010)

[4] D. W. Ming, P. D. Archer Jr., D. P. Glavin, J. L. Eigenbrode, H. B. Franz, B. Sutter, A. E. Brunner, J. C. Stern, C. Freissinet, A. C. McAdam, P. R. Mahaffy, M. Cabane, P. Coll, J. L. Campbell, S. K. Atreya, P. B. Niles, J. F. Bell III, D. L. Bish, W. B. Brinckerhoff, A. Buch, P. G. Conrad, D. J. Des Marais, B. L. Ehlmann, A. G. Fairén, K. Farley, G. J. Flesch, P. Francois, R. Gellert, J. A. Grant, J. P. Grotzinger, S. Gupta, K. E. Herkenhoff, J. A. Hurowitz, L. A. Leshin, K. W. Lewis, S. M. McLennan, K. E. Miller, J. Moersch, R. V. Morris, R. Navarro-González, A. A. Pavlov, G. M. Perrett, I. Pradler, S. W. Squyres, R. E. Summons, A. Steele, E. M. Stolper, D. Y. Sumner, C. Szopa, S. Teinturier, M. G. Trainer, A. H. Treiman, D. T. Vaniman, A. R. Vasavada, C. R. Webster, J. J. Wray, R. A. Yingst, MSL Science Team, Volatile and organic compositions of sedimentary rocks in Yellowknife Bay, Gale crater, Mars. Science 343, 1245267 (2014).

[5] V. Verma, C. Leger SSim: NASA Mars Rover Surface Flight Software Simulation. To be published in the proceedings of IEEE Aerospace Conference (2019).

Biography

__Vandi Verma__ is a Robotics Technologist and Supervisor of the Operable Robotics group in the Mobility and Robotic Systems Section at NASA JPL. She has been the SSim lead on MSL since prior to landing. She designed and developed SSim, including programming the cached sample capability. She had additional roles on MSL for the design and development of AEGIS autonomous targeting Flight Software, MSL Sample Processing Systems engineering, Flight Software Integrated Testing, Systems Integrated Testing, and Motor Control Testing. Since 2008 she has been driving Mars rovers (MER, MSL) and operating the robotic arm and sampling system. She is currently working on Autonomous Target Selection, and Robotic Arm and Sample Caching Algorithms and Flight Software for the Mars 2020 mission. She has worked on a number of research robotics projects and deployed robots in the Arctic, Antarctica and Atacama. She has a Ph.D. in Robotics from Carnegie Mellon University. Her thesis was on particle filters for robot fault detection and identification.

__Stephen Kuhn__ is a systems engineer in the Spacecraft Systems Engineering Section at NASA JPL. He received a B.S. in Mechanical Engineering from Carnegie Mellon University in 2008 and an M.S. in Aerospace Engineering from the Massachusetts Institute of Technology in 2010. He has been the flight systems engineer for the CHIMRA sampling mechanism on the Mars Science Laboratory mission, a rover planner on MSL, and held other systems engineering roles on MSL and the Mars 2020 missions at JPL.